arXiv:1402.5728v1 [q-bio.QM] 24 Feb 2014

Machine Learning Methods in the Computational Biology of Cancer

M. Vidyasagar *

February 25, 2014

Abstract

The objectives of this "perspective" paper are to review some recent advances in sparse feature selection for regression and classification, as well as compressed sensing, and to discuss how these might be used to develop tools to advance personalized cancer therapy. As an illustration of the possibilities, a new algorithm for sparse regression is presented, and is applied to predict the time to tumor recurrence in ovarian cancer. A new algorithm for sparse feature selection in classification problems is presented, and its validation in endometrial cancer is briefly discussed. Some open problems are also presented.

1 Introduction

The objectives of this "perspective" paper are to review some recent advances in sparse feature selection for regression and classification, and to discuss how these might be used in the computational biology of cancer.

Cancer is the second leading cause of death in the United States (SEER (2013)). It is estimated that in the USA in 2013, there will be 1,660,290 new cases of cancer in all sites, and 589,350 deaths (Siegel *et al.* (2013)). In the UK, in 2011 there were 331,487 cases of cancer, and 159,178 deaths. both are the latest figures available (Cancer Research UK (2013)). Worldwide, cancer led to about 7.6 million deaths in 2008 (World Health Organization

*Erik Jonsson School of Engineering and Computer Sciences University of Texas at Dallas 800 W. Campbell Road, Richardson, TX 75080, U.S.A. M.Vidyasagar@utdallas.edu.
This research was supported by National Science Foundation Awards #ECCS-1001643 and #ECCS-1306630, the Cecil & Ida Green Endowment at the UT Dallas, and by a Developmental Award from the Harold Simmons Comprehensive Cancer Center, UT Southwestern Medical Center, Dallas.

(2013)). It is interesting to note that, whether in developed countries such as the USA and the UK, or worldwide, cancer accounts for roughly 13% of all deaths (World Health Organization (2013)).

One of the major challenges faced by cancer researchers is that no two manifestations of cancer are alike, even when they occur in the same site. One can paraphrase the opening sentence of Leo Tolstoy's *Anna Karenina* and say that "Normal cells are all alike. Every malignant cell is malignant in its own way." Therefore when it comes to tackling cancer, it is essential to group these multiple manifestations into classes that are broadly similar within each class and substantially different between classes. Then attempts can be made to develop therapeutic regimens that are tailored for each class. Though this approach is often referred to in the literature as "personal" or "personalized" medicine, such nomenclature would be optimistic. It would be more accurate to describe this approach as "patient stratification." We are quite far away from truly personalized therapy at the level of a single individual. However, patient stratification is well within reach.

Until recently, grouping of cancers has been attempted first through the site of the cancer, and then through histological considerations, that is, the physical appearance of the tumor, and other parameters that can be measured by physical examination of the tumor. During the past decade, attempts have been made to collect the experimental data generated by various research laboratories into central repositories such as the Gene Expression Omnibus (GEO (2013)) and the Catalogue of Somatic Mutations in Cancer COSMIC (COSMIC (2013)). However, the data in these repositories is often collected under widely varying experimental conditions. To mitigate this problem, there are now some massive public projects under way for generating vast amounts of data for all the tumors that are available in various tumor banks, using standardized sets of experimental protocols. Amongst the most ambitious are The Cancer Genome Atlas, usually referred to by the acronym TCGA (The Cancer Genome Atlas (2013)) which is undertaken by the National Cancer Institute (NCI), and the International Cancer Genome Consortium, referred to also as ICGC (International Cancer Genomics Consortium (2013)), which is a multi-country effort. In the TCGA data, molecular measurements are available for almost all tumors, and clinical annotations are also available for many tumors. With such a wealth of data becoming freely available, researchers in the machine learning community can now aspire to make useful contributions to cancer biology without the need to undertake any experimentation themselves. Of course, without close interactions with one or more biologists, such work would be in a vacuum and have little impact.

One of the motivations for writing this paper is to present a broad picture of some recent advances in machine learning to the more mathematically inclined within the cancer biologist community, and to apply some of these techniques to a couple of problems. Full expositions of these applications will be presented elsewhere. In the other direction, it is hoped that the paper will also facilitate the entry of interested researchers from the machine learning community into cancer biology. In order to understand the *computational* aspects of the problems described here, a basic grasp of molecular biology is sufficient, as can be obtained from standard references, for example Northrop and Connor (2009); Tözeren and Byers (2003).

Now we briefly state the class of problems under discussion in this paper. This also serves to define the notation used throughout. Let m denote the number of tumor samples that are analyzed, and let n denote the number of attributes, referred to as "features," that are measured on each sample. Typically m is of the order of a few dozen in small studies, ranging up to several hundreds for large studies such as the TCGA studies, while n is of the order of tens of thousands. There are 20,000 or so genes in the human body, and in whole genome studies, and the expression level of each gene is measured by at least one "probe," and sometimes more than one. The "raw" expression level of a gene corresponds to the amount of messenger RNA (mRNA) that is produced, and is therefore a nonnegative number. However, the raw value is often transformed by taking the logarithm after dividing by a reference value, subtracting a median value, dividing by a scaling constant, and the like. As a result the numbers that are reported as gene expression levels can sometimes be negative numbers. Therefore it is best to think of gene expression levels as real numbers. Other features that are measured include micro-RNA (miRNA) levels, methylation levels, and copy number variations, all of which can be thought of as real-valued. There are also binary features such as the presence or absence of a mutation in a specific gene. In addition to these molecular attributes, there are also "labels" associated with each tumor. Let y_i denote the label of tumor i, and note that the label depends only on the sample index i and not the feature index j. Typical real-valued labels include the time of overall survival after surgery, time to tumor recurrence, or the lethality of a drug on a cancer cell line. Typical binary labels include whether a patient had metastasis (cancer spreading beyond the original site). In addition, it is also possible for labels to be ordinal variables, such as "poor responder," "medium responder," and "good responder." Often these ordinal labels are merely quantized version of some other real-valued attributes. For instance, the previous example corresponds to a three-level quantization of the time to tumor recurrence. In

general the labels refer to *clinical outcomes*, as in all of the above examples. Usually each sample has multiple labels associated with it. However, in applications, the labels are treated one at a time, so it is assumed that there is only label for each sample, with y_i denoting the label of the i-th sample. Moreover, for simplicity, it is assumed that the labels are either real-valued or binary.

Thus the measurement set can be thought of an $m \times n$ matrix $X = [x_{ij}]$, where x_{ij} is the value of feature j in sample i. The row vector x^i, denoting the i-th row of the matrix X, is called the feature vector associated with sample i. Similarly the column vector x_j denotes the variation of the j-th feature across all m samples. Throughout this paper, it is assumed that $X \in \mathbb{R}^{m \times n}$, that is, that each measurement is a real number. Binary measurements such as the presence or absence of mutations are usually handled by partitioning the data into two groups, namely those where the binary measure is zero, and where it is one. The label y_i is either bipolar (belongs to $\{-1, 1\}$) or is a real number. Taking the range of two-valued labels y_i as $\{-1, 1\}$ instead of $\{0, 1\}$ simplifies some of the formulas in the sequel. If y_i is bipoar the associated problem is called "classification" whereas if y_i is real the associated problem is called "regression." In either case, the objective is to find a function $f : \mathbb{R}^n \to \mathbb{R}$ or $f : \mathbb{R}^n \to \{-1, 1\}$ such that y_i is well-approximated by $f(x^i)$.

2 Regression Methods

The focus in this section is on the case where the label y_i is a real number. Therefore the objective is to find a function $f : \mathbb{R}^n \to \mathbb{R}$ such that $f(x^i)$ is a good approximation of y_i for all i. A typical application in cancer biology would be the prediction of the time for a tumor to recur after surgery. The data would consist of expression levels of tens of thousands of genes on around a hundred or so tumors, together with the time for the tumor to recur for each patient. The objective is to identify a small number of genes whose expression values would lead to a reliable prediction of the recurrence time. Cancer is a complex, multi-genic disease, and identifying a small set of genes that appear to be highly predictive in a particular form of cancer would be very useful. Explaining *why* these genes are the key genes would require constructing gene regulatory networks. While this problem is also amenable to treatment using statistical methods, it is beyond the scope of the present paper. Towards the end of this section, the tumor recurrence problem is studied using a new regression method.

Throughout this section, attention is focused on *linear regressors*, with $f(x) = xw - \theta$ where $w \in \mathbb{R}^n$ is a weight vector and $\theta \in \mathbb{R}$ is a threshold or bias. There are several reasons for restricting attention to linear regressors. From a mathematical standpoint, linear regressors are by far the most widely studied and the best understood class of regressors. From a biological standpoint, it makes sense to suppose that the measured outcome is a weighted linear combination of each feature, with perhaps some offset term. If one were to use higher-order polynomials for example, then biologists would rightly object that taking the product of two features (say two gene expression values) is unrealistic most of the time.[1] Other possibilities include pre-processing each feature x_{ij} through a function such as $x \mapsto e^x/(1+e^x)$, but this is still linear regression in terms of the processed values. As explained earlier, often the measured feature values x_{ij} are themselves processed values of the corresponding "raw" measurements.

In traditional least-squares regression, the objective is to choose a weight vector $w \in \mathbb{R}^n$ and a threshold θ so as to minimize the least squared error

$$J_{\rm LS} := \sum_{i=1}^{m} (x^i w - \theta - y_i)^2. \tag{1}$$

This method goes back to Legendre and Gauss, and is the staple of researchers everywhere. Let $\mathbf{e}$ denote a column vector of all ones, with the subscript denoting the dimension. Then

$$J_{\rm LS} = \|Xw - \theta \mathbf{e}_m - y\|_2^2 = \|\bar{X}\bar{w} - y\|_2^2,$$

where

$$\bar{X} = \begin{bmatrix} X & -\mathbf{e}_m \end{bmatrix} \in \mathbb{R}^{m \times (n+1)}, \bar{w} = \begin{bmatrix} w \\ \theta \end{bmatrix} \in \mathbb{R}^{n+1}.$$

If the matrix $\bar{X}$ has full column rank of $n+1$, then it is easy to see that the unique optimal choice $\bar{w}^*$ is given by

$$\bar{w}^*_{\rm LS} = (\bar{X}^t \bar{X})^{-1} \bar{X}^t y = \begin{bmatrix} X^t X & -X^t \mathbf{e}_m \\ -\mathbf{e}_m^t X & m \end{bmatrix}^{-1} \begin{bmatrix} X^t \\ \mathbf{e}_m^t \end{bmatrix} y.$$

In the present context, the fact that $m < n$ ensures that the matrix X has rank less than n, whence the matrix $\bar{X}$ has rank less than $n+1$. As a result, the standard least squares regression problem does not have a

[1] There are situations such as transcription factor genes regulating other genes, where taking such a product would be realistic. But such situations are relatively rare.

unique solution. Therefore one attempts to minimize the least-squares error while imposing various constraints (or penalties) on the weight vector w.[2] Different constraints lead to different problem formulations. An excellent and very detailed treatment of the various topics of this bsection can be found in (Hastie *et al.*, 2011, Chapter 3).

Suppose we minimize the least-squared error objective function subject to an ℓ_2-norm constraint on w. This approach to finding a unique set of weights is known as "ridge regression" and is usually credited to Hoerl and Kennard (1970). However, it would perhaps be fairer to credit the Russian mathematician A. N. Tikhonov Tikhonov (1943). In ridge regression the problem is reformulated as

$$\min \sum_{i=1}^{m} (x^i w - \theta - y_i)^2 \text{ s.t. } \|w\|_2 \leq t,$$

where t is some prespecified bound. In the associated Lagrangian formulation, the problem becomes one of minimizing the objective function

$$J_{\text{ridge}} := \sum_{i=1}^{m} (x^i w - \theta - y_i)^2 + \lambda \|w\|_2^2, \tag{2}$$

where λ is the Lagrange multiplier. Because of the additional term, the $(1,1)$-block of the Hessian of J_{ridge}, which is the Hessian of J_{ridge} with respect to w, now equals $\lambda I_n + X^t X$ which is positive definite even when $m < n$. Therefore the overall Hessian matrix is positive definite under a mild technical condition, and the problem has a unique solution for every value of the Lagrange parameter λ. However, the major disadvantage of ridge regression is that, in general, *every component* of the optimal weight vector w_{ridge} is nonzero. In the context of biological applications, this means that the regression function makes use of *every* feature x_j, which is in general undesirable.

Another possibility is to choose a solution w that has the fewest number of nonzero components, that is, a regressor that uses the fewest number of features. Define

$$\|w\|_p := \left(\sum_{i=1}^{n} |w_i|^p \right)^{1/p}$$

If $p \geq 1$, this is the familiar ℓ_p-norm. If $p < 1$, this quantity is no longer a norm, as the function $w \mapsto \|w\|_p$ is no longer convex. However, as $p \downarrow 0$,

[2] Note that no penalty is imposed on the threshold θ.

the quantity $\|w\|_p$ approaches the number of nonzero components of w. For this reason it is common to refer to the "ℓ_0-norm" even though $\|\cdot\|_0$ is not a norm at all. Moreover, it is known Natarajan (1995) that the problem of finding a $\bar{w}$ such that $\|w\|_0$ is minimized is NP-hard.

A very general formulation of the regression problem is to minimize

$$J_{\mathrm{M}} := \sum_{i=1}^{m}(x^i w - \theta - y_i)^2 + \mathcal{R}(w), \tag{3}$$

where $\mathcal{R} : \mathbb{R}^n \to \mathbb{R}_+$ is a norm known as the "regularizer." This problem is analyzed at a very high level of generality in Negabhan *et al.* (2012), where the least-squares error term is replaced by an arbitrary convex "loss" function. In the interests of simplicity, we do not discuss the results of Negabhan *et al.* (2012) in their full generality, and restrict the discussion to least-squares regression alone.

In Tibshirani (1996) it is proposed to minimize the least-squared error objective function subject to an ℓ_1-norm constraint on the weight vector w. In Lagrangian formulation, the problem is to minimize

$$J_{\mathrm{lasso}} := \sum_{i=1}^{m}(x^i w - \theta - y_i)^2 + \lambda\|w\|_1, \tag{4}$$

where λ is the Lagrange multiplier. The acronym "lasso" is coined in Tibshirani (1996), and stands for "least absolute shrinkage and selection operator". The lasso penalty can be rationalized by observing that $\|\cdot\|_1$ is the convex relaxation of the "ℓ_0-norm." The behavior of the solution to the lasso algorithm depends on the choice of the upper bound t. A detailed analysis of the Lagrangian formulation (4) and its dual problem is carried out in Osborne *et al.* (2000). It is shown there that, if the Lagrange multiplier λ in (4) is sufficiently large, say $\lambda > \lambda_{\max}$, then the only solution to the lasso minimization problem is $w = 0$. Moreover, the threshold $\lambda_{\max}$ is not easy to estimate *a priori*. An optimal solution is defined to be "regular" in (Osborne *et al.*, 2000, Definition 3.3) if it satisfies some technical conditions. In every problem there is at least one regular solution. Moreover, every regular optimal weight vector has at most m nonzero entries; see (Osborne *et al.*, 2000, Theorem 3.5).

In many applications, some of the columns of the matrix X are highly correlated. For instance, if the indices j and k correspond to two genes that are in the same biological pathway, then their expression levels would vary in tandem across all samples. Therefore the column vectors x_j and x_k would

be highly correlated. In such a case, ridge regression tends to assign nearly equal weights to each. At the other extreme, lasso tends to choose just one amongst the many correlated columns and to discard the rest; which one gets chosen is often a function of the "noise" in the measurements. In biological data sets, it is reasonable to expect that expression levels of genes that are in a common pathway are highly correlated. In such a situation, it is undesirable to choose just one amongst these genes and to discard the rest; it is also undesirable to choose all of them, as that would lead to too many features being chosen. It would be desirable to choose more than one, but not all, of the correlated columns. This is achieved by the so-called "elastic net" algorithm, introduced in Zou and Hastie (2005), which is a variation of the lasso algorithm. In this algorithm, the penalty aims to constrain, not the ℓ_1-norm of the weight w, but a weighted sum of its ℓ_1-norm and ℓ_2-norm squared. The problem formulation in this case, in Lagrangian form, is to choose w so as to minimize

$$J_{\text{EN}} := \sum_{i=1}^{n} (x^i w - \theta - y_i)^2 + \lambda[\mu \|w\|_2^2 + (1-\mu)\|w\|_1], \tag{5}$$

where $\mu \in (0,1)$. Note that if $\mu = 0$, then the elastic net algorithm becomes the lasso, whereas with $\mu = 1$, the elastic net algorithm becomes ridge regression. Thus the elastic net algorithm provides a bridge between the two. Note that the penalty term in the elastic net algorithm is *not* a norm, due to the presence of the squared term; hence the elastic net algorithm is not covered by the very thorough analysis in Negabhan *et al.* (2012). A useful property of the elastic net algorithm is brought out in (Zou and Hastie, 2005, Theorem 1).

Theorem 2.1 *Assume that y, X, λ are fixed, and let $\bar{w}$ denote the corresponding minimizer of (5). Assume without loss of generality that y is centered, that is, $y^t \mathbf{e}_m = 0$, and that the columns of X are normalized such that $\|x_j\|_2 = 1$ for all j. Let j, k be two indices between 1 and n, and suppose that $x_j^t x_k \geq 0$. Then*

$$|w_j - w_k| \leq \frac{\|y\|_1}{\lambda\mu} \sqrt{2(1 - x_j^t x_k)}. \tag{6}$$

Since one can always ensure that $x_j^t x_k^t \geq 0$ by replacing x_k by $-x_k$ if necessary, (6) states that if the columns x_j and x_k are highly correlated, then the corresponding coefficients in the regressor are nearly equal. Unlike in the lasso algorithm, there do not seem to be many results on the number

of nonzero weights that are chosen by the elastic net algorithm. It can and often does happen that the number of features chosen is larger than m, the number of samples. However, as explained above, this is often seen as a desirable feature when the columns of the matrix X are highly correlated, as they often are in biology data sets.

Next we discuss several versions of the problem formulation in (3) corresponding to diverse choices of the penalty norm $\mathcal{R}$, culminating in some open problems that are relevant to biological applications. The "pure" lasso algorithm tries to choose as few distinct features as possible in the regressor. However, it may be worthwhile to partition the set of features $\mathcal{N} = \{1, \ldots, n\}$ into g groups $G_1, \ldots, G_g$, and then choose a regressor that selects elements from as few distinct groups as possible, without worrying about the number of features chosen. This is achieved by the so-called group lasso algorithm introduced in Bakin (1999) and Lin and Zhang (2006). Let $n_l := |G_l|$ for $l = 1, \ldots, g$. In the grouped lasso algorithm, the objective function is

$$J_{\rm GL} = \sum_{i=1}^{m} (x^i w - \theta - y_i)^2 + \lambda \sum_{l=1}^{g} \sqrt{n_l} \|w_{G_l}\|_2, \tag{7}$$

where $w_{G_l} \in \mathbb{R}^n$ is determined from w by setting $w_j = 0$ for all $j \notin G_l$. It is clear that, depending on the relative sizes of the various groups, one weight vector can have more nonzero components than another, and yet the number of distinct groups to which these nonzero components belong can be smaller. In the limiting case, if the number of groups is taken as n and each group is taken to consist of a singleton set, then the grouped lasso reduces to the standard lasso algorithm. A further variation is the so-called sparse group lasso algorithm introduced in Friedman *et al.* (2010); Simon *et al.* (2012), where the objective is simultaneously to choose features from as few distinct groups as possible, and within the chosen groups, choose as few features as possible. The objective function in the sparse group lasso (SGL) algorithm is

$$J_{\rm SGL} = \sum_{i=1}^{m} (x^i w - \theta - y_i)^2 + \lambda \sum_{l=1}^{g} [(1-\mu)\|w_{G_l}\|_1 + \mu \|w_{G_l}\|_2], \tag{8}$$

where as always $\mu \in [0, 1]$.

The above formulations of the GL and SGL norms based on the assumption that the various groups do not overlap. However, in some biological applications it makes sense to permit overlapping group decompositions. Specifically, at a first-level of approximation a gene regulatory network

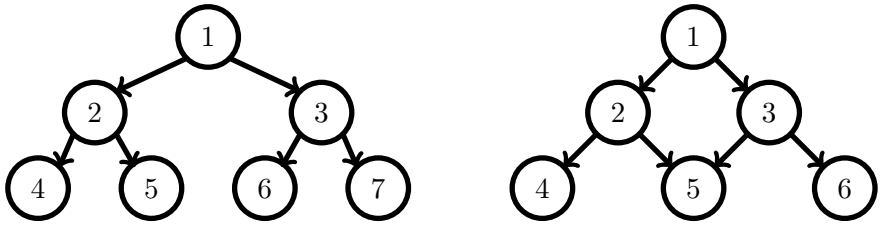

Figure a: Two regulatory networks. (i) A network without overlapping groups. (ii) A network with overlapping groups.

(GRN) can be modelled as a directed acyclic graph (DAG), wherein the root nodes can be interpreted as master regulator genes, and directed paths can be interpreted as biological pathways. In such a case, one seeks to explain the available data, not by choosing the fewest number of *genes*, but rather by the fewest number of *pathways*. To illustrate, consider the baby example shown in Figure a, where genes 1 and 2 are master regulators, while genes 3 through 7 are regulated genes. Some are regulated directly by a master regulator gene, while others are indirectly regulated. In Figure a(i), there are four pathways, namely

$$G_1 = \{1, 2, 4\}, G_2 = \{1, 2, 5\}, G_3 = \{1, 3, 6\}, G_4 = \{1, 3, 7\},$$

whereas in Figure a(ii) there are also four pathways, namely

$$G_1 = \{1, 2, 4\}, G_2 = \{1, 2, 5\}, G_3 = \{1, 3, 5\}, G_4 = \{1, 3, 6\}.$$

Ideally, we would like to choose a set of features that intersect with as few pathways as possible. We will return to this example after presenting available theories for sparse regression with overlapping groups.

To date, various versions of group or sparse group lasso with overlapping groups have been proposed. As before, let $G_1, \ldots, G_g$ be subsets of $\mathcal{N} = \{1, \ldots, n\}$, but now *without* the assumption that the groups are pairwise disjoint. The penalty-augmented optimization problems are the same as in (7) and (8) respectively; however, the objective functions are now referred to as J_{GLO} and J_{SGLO} to suggest (sparse) group lasso with overlap. For the case of overlapping groups, the theory developed in Negabhan *et al.* (2012) continues to apply so long as the penalty terms in (7) and (8) respectively are "decomposable." The most general results available to date address the case where the groups are "tree-structured," that is,

$$G_i \cap G_j \neq \emptyset \;\Rightarrow\; \{G_i \subseteq G_j \text{ or } G_j \subseteq G_i\}. \tag{9}$$

See for example Obozinski *et al.* (2011); Jenetton *et al.* (2011).

Now, if we examine the groups associated with the network in Figure a(i), it is obvious that (9) is not satisfied. However, there is a slight modification that would permit (9) to hold, namely to drop the root node and retain only the successors. Thus the various groups are

$$G_1 = \{4\}, G_2 = \{5\}, G_3 = \{6\}, G_4 = \{7\},$$

$$G_5 = \{2, 4\}, G_6 = \{2, 5\}, G_7 = \{3, 6\}, G_8 = \{3, 7\}.$$

However, there is no way of modifying the groups so as to ensure that (9) holds for the network in Figure a(ii). The reason is easy to see. The "tree structure" assumption (9) implies that there is only one path between every pair of nodes. But this is clearly not true in Figure a(ii), because there are two distinct paths from node 1 to node 5. Moreover, a little thought would reveal that that the assumption of tree-structured groups does not really permit truly overlapping groups. In particular, if (9) holds, then the collection of sets $\{G_1, \ldots, G_g\}$ can be expressed as a union of chains in the form

$$G_{11} \subseteq \ldots \subseteq G_{1g_1}, \ldots, G_{s1} \subseteq \ldots \subseteq G_{sg_s},$$

where the "maximal" sets G_{ig_i} are pairwise disjoint and together span the total feature set $\mathcal{N} = \{1, \ldots, n\}$. Now, in a biological network, it makes no sense to impose a condition that there must be only path between every pair of nodes. Therefore the problem of defining a decomposable norm penalty for inducing other types of sparsity besides tree-structure, especially the types of sparsity that are consistent with biology, is still open.

We conclude this section with a new algorithm and its application to sparse regression. This represents joint work with Mehmet Eren Ahsen and will be presented in more complete form elsewhere. A special case of SGL is obtained by choosing just one group, which perforce has to equal $\mathcal{N}$, so that

$$J_{\text{MEN}} = \sum_{i=1}^{m} (x^i w - \theta - y_i)^2 + \lambda[(1-\mu)\|w\|_1 + \mu\|w\|_2]. \tag{10}$$

Of course, since the entire index set $\mathcal{N}$ is chosen as one group, there is nothing "sparse" about it. Note that the only difference between (10) and (5) is that the ℓ_2-norm is *not* squared in the former. For this reason, the above approach is called the " modified elastic net" or MEN algorithm. Unlike in EN, the penalty (or constraint) term in MEN is a norm, being a convex combination of the ℓ_1- and ℓ_2-norms. In several examples, the MEN algorithm appears to combine the accuracy of EN with the sparsity

of lasso. It is relatively easy to prove an analog of Theorem 2.1 for the MEN algorithm. That is, unlike in lasso but as in EN, MEN assigns nearly equal weights to highly correlated features. But further theoretical analysis remains to be carried out.

The MEN algorithm was applied to the TCGA ovarian cancer data (The Cancer Genome Atlas Network (2011)) to predict the time to tumor recurrence. Specifically, both times to tumor recurrence as well as expression levels for 12,042 genes are available for 283 patients. Out of these, 40 patients whose tumors recurred before 210 days or after 1,095 days were excluded from the study as being "extreme" cases. The remaining 243 samples were analyzed using MEN with recursive feature elimination. The results are shown in Figure b. The number of features and the average percentage error in absolute value are shown in Table 1.

Algorithm	**No. of Features**	**Average Perc. Error**
LASSO	43	16.14%
Elastic Net	60	14.35&
MEN	42	14.91%

Table 1: Comparison of three algorithms on TCGA ovarian cancer data on time to tumor recurrence, with extreme cases excluded.

3 Compressed Sensing

In recent years, there have been several results that are grouped under the general heading of "compressed sensing" or "compressive sensing". Both expressions are in use, but "compressed sensing" is used in this paper. The problem can be roughly stated as follows: Suppose $x \in \mathbb{R}^n$ is an unknown vector but with known structure; is it possible to determine x either exactly or approximately, by taking $m \ll n$ linear measurements of x? The area of research that goes under this broad heading grew spectacularly during the first decade of the new millennium.[3] As summarized in the introduction of the paper Donoho (2006a), the impetus for recent work in this area was the desire to find algorithms for data compression that are "universal" in the sense of being nonadaptive (i.e., do not depend on the data). In

[3]In Davenport *et al.* (2012) it is suggested a precursor of compressed sensing can be found in a paper that dates back to 1795!

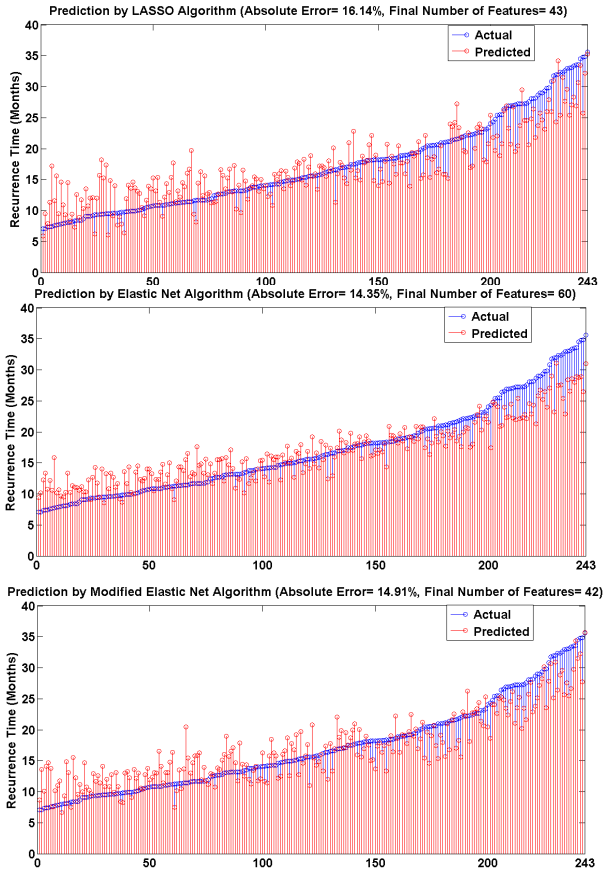
Prediction by LASSO Algorithm (Absolute Error= 16.14%, Final Number of Features= 43)
Actual
Predicted
Recurrence Time (Months)
Prediction by Elastic Net Algorithm (Absolute Error= 14.35%, Final Number of Features= 60)
Actual
Predicted
Recurrence Time (Months)
Prediction by Modified Elastic Net Algorithm (Absolute Error= 14.91%, Final Number of Features= 42)
Actual
Predicted
Recurrence Time (Months)

Figure b: Predicted vs. actual times to tumor recurrence in 243 ovarian cancer patients. The results for lasso are at the top, those for the elastic net are in the middle and those for the modified elastic net algorithm are at the bottom.

the original papers in this area, the results and proofs were a mixture of sampling, signal transformation (time domain to frequency domain and vice versa), randomness etc. However, as time went on, the essential ingredients of the approach were identified, thus leading to a very streamlined theory that clearly transcends its original application domains of image and signal processing.

The motivation for discussing compressed sensing theory in the present paper is the following: Whether it is in compressed sensing or in computational biology, one searches for a relatively simple explanation of the observations. Therefore it may *potentially* be possible to borrow some of the basic ideas from compressed sensing theory and adapt them to problems in cancer biology. Compressed sensing theory *as it currently stands* cannot directly be applied to the analysis of biological data sets, because the fundamental assumption in compressed sensing theory is that *one is able to choose* the so-called measurement matrix, called A throughout this paper. Note that in statistics, the matrix A is often referred to as the "design" matrix. However, in biological applications this matrix is often fixed, and one does not have the freedom to choose, i.e. to "design" it. However, in biological (and other) applications, the measurement matrix is given, and one does not have the freedom to change it. Nevertheless, the developments in this area are too important to be ignored by computational biologists. The hope is that, by understanding the core arguments of compressed sensing theory and building on them, it will be possible for the computational biology community to develop a similarly successful theory for their application domain. Therefore an introductory treatment of compressed sensing is included here.

The major developments in this area are generally associated with the names of Candès, Donoho, Romberg, and Tao, though several other researchers have also made important contributions. See Donoho (2006a) for one of the earliest comprehensive papers, as well Donoho (2006a,b); Candès (2008); Candès and Tao (2005, 2007); Candès and Plan (2009); Romberg (2009); Cohen *et al.* (2009). The survey paper Davenport *et al.* (2012) and a recent paper Negabhan *et al.* (2012) contain a wealth of bibliographic references that can be followed up by interested readers.

We begin by introducing some notation. Suppose m, n, k are given integers, with $n \geq 2k$. For convenience, we denote the set $\{1, \ldots, n\}$ by $\mathcal{N}$ throughout. For a given vector $x \in \mathbb{R}^n$, let $\text{supp}(x)$ denote its support, that is, $\text{supp}(x) = \{i : x_i \neq 0\}$. Let $\Sigma_k = \{x \in \mathbb{R}^n : |\text{supp}(x)| \leq k\}$. Thus Σ_k denotes the set of "k-sparse" vectors in $\mathbb{R}^n$, or in other words, the set of n-dimensional vectors that have k or fewer nonzero components. For each vector $x \in \mathbb{R}^n$, integer $k < n$, and norm $\|\cdot\|$ on $\mathbb{R}^n$, the symbol $\sigma_k(x, \|\cdot\|)$

denotes the distance from x to Σ_k, that is,

$$\sigma_k(x, \|\cdot\|) = \inf\{\|x - z\| : z \in \Sigma_k\}.$$

The quantity $\sigma_k(x, \|\cdot\|)$ is called the "sparsity measure" of the vector x of order k with respect to the norm $\|\cdot\|$. It is obvious that $\sigma_k(x, \|\cdot\|)$ depends on the underlying norm. However, if $\|\cdot\|$ is one of the ℓ_p-norms, then it is easy to compute $\sigma_k(x, \|\cdot\|)$. Specifically, given k, let Λ_0 denote the index set corresponding to the k-largest components of x in magnitude, and let $x_{\Lambda_0^c}$ denote the vector that results by replacing the components of x in the set Λ_0 by zeros. (It is convenient to think of x_{Λ^c} as an element of $\mathbb{R}^n$ rather than an element of $\mathbb{R}^{n-k}$.) Then, whenever $p \in [1, \infty]$, it is easy to see that

$$\sigma_k(x, \|\cdot\|_p) = \|x_{\Lambda_0^c}\|_p.$$

Next, the so-called " restricted isometry property" (RIP) is introduced. Note that in some cases the RIP can be replaced by a weaker property known as the "null space property" (Cohen *et al.* (2009)). However, the objective of the present paper is *not* to present the most general results, but rather to present reasonably general results that are easy to explain. So the exposition below is confined to the RIP.

Definition 3.1 *Suppose $A \in \mathbb{R}^{m\times n}$. We say that A satisfies the* **Restricted Isometry Property (RIP)** *of order k with constant δ_k if*

$$(1 - \delta_k)\|u\|_2^2 \leq \langle u, Au\rangle \leq (1 + \delta_k)\|u\|_2^2, \ \forall u \in \Sigma_k. \tag{11}$$

So the matrix A has the RIP of order k with constant $1 - \delta_k$ if the following property holds: For every choice of k or fewer columns of A (say the columns in the set $J \subseteq \mathcal{N}$ where $|J| \leq k$), the spectrum of the symmetric matrix $A_J^t A_J$ lies in the interval $[1 - \delta_k, 1 + \delta_k]$, where $A_J \in \mathbb{R}^{m\times|J|}$ denotes the submatrix of A consisting of all rows and the columns corresponding to the indices in J.

If integers n, k are specified, the integer m has to be sufficiently large in order for the matrix A to satisfy the RIP.

Theorem 3.1 *((Davenport* et al.*, 2012, Theorem 1.4)) Suppose $A \in \mathbb{R}^{m\times n}$ satisfies the RIP or order $2k$ with constant $\delta_{2k} \in (0, 1/2]$. Then*

$$m \geq ck \log(n/k) = ck(\log n - \log k), \tag{12}$$

where

$$c = \frac{1}{2\log(\sqrt{24} + 1)} \approx 0.28.$$

Next we state some of the main known results in compressed sensing. The theorem statement below corresponds to (Candès, 2008, Theorem 1.2) and (Davenport *et al.*, 2012, Theorem 1.9). It can be compared with (Candès and Plan, 2009, Theorem 1.4)).

Theorem 3.2 *Suppose $A \in \mathbb{R}^{m\times n}$ satisfies the RIP of order δ_{2k} with constant $\delta_{2k} < \sqrt{2}-1$, and that $y = Ax + \eta$ for some $x \in \mathbb{R}^n$ and $\eta \in \mathbb{R}^m$ with $\|\eta\|_2 \leq \epsilon$. Let $\mathcal{B}_\epsilon(y) := \{z \in \mathbb{R}^n : \|y - Az\|_2 \leq \epsilon\}$, and define*

$$\hat{x} = \underset{z\in\mathcal{B}_\epsilon(y)}{\operatorname{argmin}} \|z\|_1. \tag{13}$$

Then

$$\|\hat{x} - x\|_2 \leq C_0 \frac{\sigma_k(x, \|\cdot\|_1)}{\sqrt{k}} + C_2\epsilon, \tag{14}$$

where

$$C_0 = 2\frac{1 + (\sqrt{2}-1)\delta_{2k}}{1 - (\sqrt{2}+1)\delta_{2k}}, C_2 = \frac{4\sqrt{1+\delta_{2k}}}{1 - (\sqrt{2}+1)\delta_{2k}}. \tag{15}$$

The formula for C_2 is written slightly differently from that in (Davenport *et al.*, 2012, Theorem 1.9) but is equivalent to it.

Corollary 3.1 *Suppose $A \in \mathbb{R}^{m\times n}$ satisfies the RIP of order δ_{2k} with constant $\delta_{2k} < \sqrt{2}-1$, and that $y = Ax + \eta$ for some $x \in \Sigma_k$ and $\eta \in \mathbb{R}^m$ with $\|\eta\|_2 \leq \epsilon$. Let $\mathcal{B}_\epsilon(y) := \{z \in \mathbb{R}^n : \|y - Az\|_2 \leq \epsilon\}$, and define*

$$\hat{x} = \underset{z\in\mathcal{B}_\epsilon(y)}{\operatorname{argmin}} \|z\|_1. \tag{16}$$

Then

$$\|\hat{x} - x\|_2 \leq C_2\epsilon, \tag{17}$$

where C_2 is defined in (15).

Corollary 3.2 *Suppose $A \in \mathbb{R}^{m\times n}$ satisfies the RIP of order δ_{2k} with constant $\delta_{2k} < \sqrt{2}-1$, and that $y = Ax$ for some $x \in \Sigma_k$. Let $A^{-1}(y) := \{z \in \mathbb{R}^n : y = Az\}$, and define*

$$\hat{x} = \underset{z\in A^{-1}(y)}{\operatorname{argmin}} \|z\|_1. \tag{18}$$

Then $\hat{x} = x$.

Both corollaries follow readily from the bound (14). Note that if $x \in \Sigma_k$ then $\sigma_k(x, \|\cdot\|_1) = 0$. Thus (14) implies that $\|\hat{x} - x\|_2 \leq C_2\epsilon$ if there is measurement error, and $\|\hat{x} - x\|_2 = 0$, i.e. that $\hat{x} = x$, if there no measurement error.

Corollary 3.1 is a simplified version of a corresponding result in Candès and Plan (2009) and is referred to as the "near ideal" property of the lasso algorithm. Suppose that $x \in \Sigma_k$ so that x is k-sparse. Let S denote the support of x, and let $A_S \in \mathbb{R}^{m \times |S|}$ denote the submatrix of A consisting of the columns corresponding to indices in S. If an "oracle" knew not only the size of S, but the set S itself, then the oracle could compute $\hat{x}$ as

$$\hat{x}_{\text{oracle}} = (A_S^T A_S)^{-1} A_S^T y = x + (A_S^T A_S)^{-1} A_S^T \eta.$$

Then the error would be

$$\|\hat{x}_{\text{oracle}} - x\|_2 = \|(A_S^T A_S)^{-1} A_S^T \eta\|_2 \leq \text{const} \cdot \epsilon$$

for some appropriate constant. On the other hand, if $x \in \Sigma_k$, then $\sigma_k(x, \|\cdot\|_1) = 0$, and the right side of (14) reduces to (17), that is,

$$\|\hat{x} - x\|_2 \leq C_2\epsilon.$$

The point therefore is that, if the matrix A satisfies RIP, and the constant δ_{2k} satisfies the "compressibility condition" $\delta_{2k} < \sqrt{2} - 1$, then the mean-squared error of the solution to the optimization problem (16) is bounded by a fixed (or "universal") constant times the error bound achieved by an "oracle" that knows the support of x.

The advantage of the above theorem statements, which are taken from Candès (2008); Davenport *et al.* (2012), is that the role of various conditions is clearly delineated. For instance, the construction of a matrix $A \in \mathbb{R}^{m \times n}$ that satisfies the RIP is usually achieved by some randomized algorithm. In (Candès and Tao, 2005, Theorem 1.5) such a matrix is constructed by taking the columns of A to be samples of i.i.d. Gaussian variables. In Achlioptas (2003), Bernoulli processes are used to construct A, which has the advantage of ensuring that all elements a_{ij} have just three possible values, namely $0, +1, -1$. A simple proof that the resulting matrices satisfy the RIP with high probability is given in Baraniuk *et al.* (2008). Neither of these construction methods is *guaranteed* to generate a matrix A that satisfies RIP. Rather, the resulting matrix A satisfies RIP with some probability, say $\geq 1 - \gamma_1$. The probability γ_1 that the randomized method may fail to generate a suitable A matrix can be bounded using techniques that have nothing to do with

the above theorem. Similarly, in case the measurement matrix A satisfies the RIP but the measurement noise η is random, then it is obvious that Theorem 3.2 holds with probability $\geq 1 - \gamma_2$, where γ_2 is a bound on the tail probability $\Pr\{\|\eta\|_2 > \epsilon\}$. Again, the problem of bounding this tail probability has nothing to do with Theorem 3.2. By combining both estimates, it follows that if the measurement matrix A is generated through randomization, and if the measurement noise is also random, then Theorem 3.2 holds with probability $\geq 1 - \gamma_1 - \gamma_2$.

Observe that the optimization problem (16) is

$$\min_z \|z\|_1 \text{ s.t. } \|y - Az\|_2 \leq \epsilon.$$

This is essentially the same as the lasso algorithm, except that the objective function and the constraint are interchanged. This raises the question as to whether the ℓ_1-norm can be replaced by some other norm $\|\cdot\|_P$ that induces some other form of sparsity, for example group sparsity. If some other norm is used in place of the ℓ_1-norm, does the resulting algorithm display near-ideal behavior, as does lasso? In other words, is there an analog of Theorem 3.2 if $\|\cdot\|_1$ is replaced by another penalty $\|\cdot\|_P$? In joint work with Mehmet Eren Ahsen (Ahsen and Vidyasagar (2014)), the author has proved a very general theorem to the following effect: Whenever the penalty norm is "decomposable" and the measurement matrix A satisfies a "group RIP," the corresponding algorithm has near-ideal behavior provided a "compressibility condition" is satisfied. The result is described in brief.

Let $\mathcal{G} = \{G_1, \ldots, G_g\}$ be a partition of $\mathcal{N} = \{1, \ldots, n\}$. This implies that the sets G_i are pairwise disjoint. If $S \subseteq \{1, \ldots, g\}$, define $G_S := \cup_{i \in S} G_i$. Let k be some integer such that $k \geq \max_i |G_i|$. A subset $\Lambda \subseteq \mathcal{N}$ is said to be S**-group** k**-sparse** if $\Lambda \subseteq G_S$ and $|G_S| \leq k$, and **group** k**-sparse** if it is S-group k-sparse for some set $S \subseteq \{1, \ldots, g\}$. The symbol GkS $\subseteq 2^{\mathcal{N}}$ denotes the collection of group k-sparse sets.

Suppose $\|\cdot\|_P : \mathbb{R}^n \to \mathbb{R}_+$ is some norm. The next definition builds on an earlier definition from Negabhan *et al.* (2012).

Definition 3.2 $\|\cdot\|_P$ *is* **decomposable** *with respect to the partition $\mathcal{G}$ if the following is true: Whenever $u, v \in \mathbb{R}^n$ are group k-sparse with support sets $\Lambda_u \subseteq G_{S_1}$, $\Lambda_v \subseteq G_{S_2}$ and the sets S_1, S_2 are disjoint, it is true that*

$$\|u + v\|_P = \|u\|_P + \|v\|_P. \tag{19}$$

The norm $\|\cdot\|_P$ is said to be **regular** *if, whenever Λ is a strict subset of G_i for some index i, it is true that $\|x_\Lambda\|_P \leq \|x_G\|_P$ with equality if and only if $x_{G \setminus \Lambda} = 0$.*

By adapting the arguments in Negabhan *et al.* (2012), it can be shown that the group lasso norm used in (7), namely

$$\|x\|_{\mathrm{GL}} := \sum_{l=1}^{g} \sqrt{n_l} \|x_{G_l}\|_2,$$

and the sparse group lasso norm used in (8)

$$\|x\|_{\mathrm{SL}} := \sum_{l=1}^{g} [(1-\mu)\|x_{G_l}\|_1 + \mu \|x_{G_l}\|_2],$$

are both decomposable as well as regular.

Next, the notion of RIP is extended to groups.

Definition 3.3 *A matrix $A \in \mathbb{R}^{m\times n}$ is said to* **satisfy the group RIP of order k with constants $\underline{\rho}_k, \bar{\rho}_k$** *if*

$$0 < \underline{\rho}_k \leq \min_{\Lambda \in \mathrm{GkS}} \min_{supp(z)\subseteq\Lambda} \frac{\|Az\|_2^2}{\|z\|_2^2} \leq \max_{\Lambda \in \mathrm{GkS}} \max_{supp(z)\subseteq\Lambda} \frac{\|Az\|_2^2}{\|z\|_2^2} \leq \bar{\rho}_k. \tag{20}$$

We define $\delta_k := (\bar{\rho}_k - \underline{\rho}_k)/2$, and introduce some constants:

$$c := \min_{\Lambda \in \mathrm{GkS}} \min_{x_\Lambda \neq 0} \frac{\|x_\Lambda\|_P}{\|x_\Lambda\|_2}, d := \max_{\Lambda \in \mathrm{GkS}} \max_{x_\Lambda \neq 0} \frac{\|x_\Lambda\|_P}{\|x_\Lambda\|_2}. \tag{21}$$

With these definitions, the following theorem can be proved.

Theorem 3.3 *Suppose $A \in \mathbb{R}^{m\times n}$ satisfies the group RIP property of order $2k$ with constants $(\underline{\rho}_{2k}, \bar{\rho}_{2k})$ respectively, and let $\delta_{2k} = (\bar{\rho}_{2k} - \underline{\rho}_{2k})/2$. Suppose $x \in \mathbb{R}^n$ and that $y = Ax + \eta$ where $\|\eta\|_2 \leq \epsilon$. Suppose that the norm $\|\cdot\|_P$ is decomposable, and define*

$$\hat{x} = \operatorname*{argmin}_{z\in\mathbb{R}^n} \|z\|_P \ \ s.t. \ \ \|y - Az\|_2 \leq \epsilon. \tag{22}$$

Suppose that

$$\delta_{2k} < \frac{c\underline{\rho}_k}{d}. \tag{23}$$

Then

$$\|\hat{x} - x\|_P \leq \frac{2}{1-r}[2\sigma + (1+r)\zeta\epsilon], \tag{24}$$

and

$$\|\hat{x} - x\|_2 \leq \frac{2}{c(1-r)}[2\sigma + (1+r)\zeta\epsilon], \tag{25}$$

where

$$r := \frac{\delta_{2k} d}{c\underline{\rho}_k}, \sigma := \sigma_{k,\mathcal{G}}(x, \|\cdot\|_P), \zeta := \frac{2d\sqrt{\bar{\rho}_k}}{\underline{\rho}_k}, \tag{26}$$

and c, d are defined in (21).

In the above theorem, (23) replaces the "compressibility" condition $\delta_{2k} < \sqrt{2} - 1$ of Theorem 3.2. The resemblance of (25) to (14) is obvious. Consequently, (25) can be readily interpreted as stating that minimizing the decomposable norm $\|\cdot\|\|_P$ leads to near-ideal behavior.

4 Classification Methods

The basic problem of classification can be stated as follows: Suppose we are given vectors $x^i, i = 1, \ldots, m$ where each $x^i \in \mathbb{R}^n$ is viewed as a row vector. Suppose further that the m vectors are grouped into two sets, call them $\mathcal{M}_1$ and $\mathcal{M}_2$. Without loss of generality, renumber the vectors such that $x_1, \ldots, x_{m_1}$ constitute the first set and $x_{m_1+1}, \ldots, x_{m_1+m_2} = x_m$ constitute the second set. For future use, define $\mathcal{M} = \{1, \ldots, m\}$, and let $\mathcal{M}_1 = \{1, \ldots, m_1\}$, $\mathcal{M}_2 = \{m_1+1, \ldots, m_1+m_2 = m\}$ be a partition of $\mathcal{M}$. Assign a label $y_i = +1$ to the vectors in $\mathcal{M}_1$ and a label $y_i = -1$ to the vectors in $\mathcal{M}_2$. The objective of (two-class) classification is to find a **discriminant function** $f : \mathbb{R}^n \to \mathbb{R}$ such that $f(x^i)$ has the same sign as y_i for all i, or equivalently $y_i \cdot \text{sign}(f(x^i)) = 1$ for all i. In the present context, the objective is not merely to find such a discriminant function, but rather, to find one that uses relatively few features.

In many ways, classification is an easier problem than regression, because the sole criterion is that the discriminant function $f(x^i)$ should have the same sign as the label y_i for each i. Thus, if f is a discriminant function, so is αf for every positive constant α, and more generally, so is any function $\phi(f)$ whenever ϕ is a so-called "first and third-quadrant function," i.e. where $\phi(u) > 0$ when $u > 0$ and $\phi(u) < 0$ when $u < 0$. This gives us great latitude in choosing a discriminant function.

4.1 The Support Vector Machine

This section is devoted to the well-known Support Vector Machine (SVM), first introduced in Cortes and Vapnik (1997), which is amongst the most successful and most widely used tools in machine learning.

A given set of labelled vectors $\{(x^i, y_i), x^i \in \mathbb{R}^n, y_i \in \{-1, 1\}\}$ is said to be **linearly separable** if there exist a "weight vector" $w \in \mathbb{R}^n$ (viewed as

a column vector) and a "threshold" $\theta \in \mathbb{R}$ such that $f(x) = xw - \theta$ serves as a discriminant function. Equivalently, the data set is linearly separable if there exist a weight vector $w \in \mathbb{R}^n$ and a threshold $\theta \in \mathbb{R}$ such that

$$x^i w > \theta \; \forall i \in \mathcal{M}_1, x^i w < \theta \; \forall i \in \mathcal{M}_2.$$

To put it yet another way, given a weight w and a threshold θ, define $\mathcal{H} = \mathcal{H}(w, \theta)$ by

$$\mathcal{H} := \{x \in \mathbb{R}^n : xw - \theta = 0\}. \mathcal{H}_+ := \{x \in \mathbb{R}^n : xw - \theta > 0\}, \mathcal{H}_- := \{x \in \mathbb{R}^n : xw - \theta < 0\}.$$

The data set is linearly separable if there exists a hyperplane $\mathcal{H}$ such that $x^i \in \mathcal{H}_+ \; \forall i \in \mathcal{M}_1$, and $x^i \in \mathcal{H}_- \; \forall i \in \mathcal{M}_2$.

The situation can be depicted as in Figure c, where the vermilion dots denote vectors in Class $\mathcal{M}_1$ and the dark green dots denote vectors in Class $\mathcal{M}_2$; however, it is clear that linear separability is not affected by swapping the class labels.

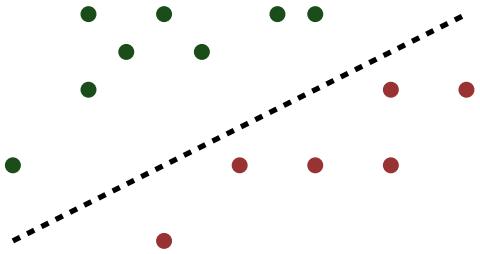

Figure c: A Linearly Separable Data Set

It is easy to determine whether or not a given data set is linearly separable, because that is equivalent to the feasibility of a linear programming problem. More general results can be obtained using Vapnik-Chervonenkis theory; see Wenocur and Dudley (1981); Vidyasagar (2003). Suppose that the n vectors $x_1, \ldots, x_n$ do not lie on an $(p-1)$-dimensional hyperplane in $\mathbb{R}^p$. In such a case, whenever $p \geq n - 1$, the data set is linearly separable for *every one* of the 2^n ways of assigning labels to the n vectors. This result suggests that, if a given data set is not linearly separable, it can be made so by increasing the dimension of the data vectors x^i, for instance, by including not just the original components but also their higher powers. This is the rationale behind so-called "higher-order" SVMs, or more generally, kernel-based classifiers; see e.g. Cristianini and Shawe-Taylor (2000); Schölkopf and Smola (2002).

It is also easy to see that, if there exists *one* hyperplane that separates the two classes, there exist *infinitely many* such hyperplanes. The Support Vector Machine (SVM) introduced in Cortes and Vapnik (1997) chooses the separating hyperplane such that *the nearest point to the hyperplane within each class is as far as possible from it.* In the original SVM formulation, the distance to the hyperplane is measured using the Euclidean or ℓ_2-norm. To illustrate the concept, the same data set as in Figure c is shown again in Figure d, with the "optimal" separating hyperplane, and the closest points to it within the two classes shown as hollow circles.

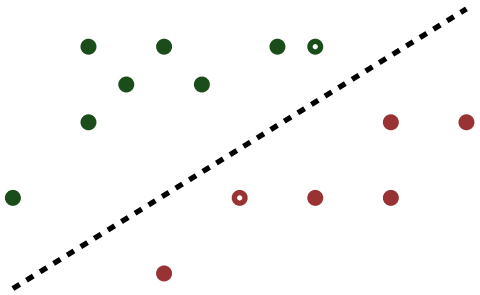

Figure d: Optimal Separating Hyperplane

In symbols, the SVM is obtained by solving the following optimization problem:

$$\max_{w,\theta} \min_{i} \inf_{v \in \mathcal{H}} \|v - x^i\|.$$

An equivalent formulation of the SVM is obtained by observing that the distance of the separating hyperplane to the nearest points is given by $c/\|w\|$, where

$$c := \min_{i \in \mathcal{M}_1} |y_i(x^i w - \theta)| = \min_{i \in \mathcal{M}_2} |y_i(x^i w - \theta)|,$$

where the equality of the two terms follows from the manner in which the separating hyperplane is chosen. Moreover, the optimal hyperplane is invariant under scale change, that is, multiplying w and θ by a positive constant. Therefore there is no loss of generality in taking the constant c to equal one. With this rescaling, the problem at hand becomes the following:

$$\min_{w} \|w\| \text{ s.t. } x^i w \geq 1 \; \forall i \in \mathcal{M}_1, x^i w \leq -1 \; \forall i \in \mathcal{M}_2. \tag{27}$$

This is the manner in which the SVM is implemented nowadays in most software packages.

If the norm in (27) is the ℓ_2-norm, then the minimization problem (27) is a quadratic programming problem, which can be solved efficiently for extremely large data sets. Moreover, the introduction of new data points does not alter the optimal hyperplane, unless one of the new data points is closer to the hyperplane than the earlier closest points. This is illustrated in Figure e, which contains exactly the same vectors as in Figure d, plus two more shown in blue and red respectively. The optimal hyperplane remains the same. For all these reasons, the SVM offers a very attractive approach to finding a classifier in situations where the number of features is smaller than the number of samples.

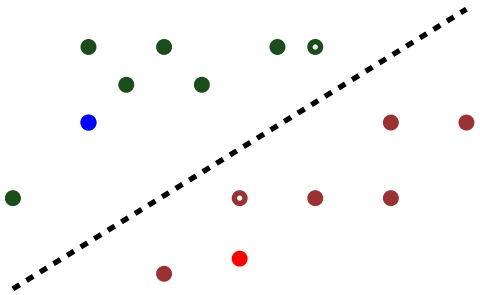

Figure e: Optimal Separating Hyperplane

Unfortunately, in biological applications, the situation is usually the reverse: The number of features (the dimension of the vectors x^i) is a few orders of magnitude larger than the number of samples (the number of vectors). In such a case, because of the results in Wenocur and Dudley (1981), linear separability is not an issue. However, in general, *every component* of the optimal weight vector w is nonzero. This means that a classifier uses every single feature in order to discriminate between the classes. Clearly this is undesirable.

4.2 The ℓ_1-Norm Support Vector Machine

In this subsection, we introduce a modification of the original SVM formulation due to Bradley and Mangasarian (1998). At the same time, we also incorporate a further feature for trading off the false positive error rate and the false negative error rate, due originally to Veropoulos *et al.* (1999).

The original SVM formulation presupposes that the data set is linearly separable. This naturally raises the question of what is to be done in case the data set is *not* linearly separable. One way to approach the problem is

to choose a hyperplane that misclassifies the fewest number of points. While appealing, this approach is impractical, because it is known that this problem is NP-hard; see Höffgen *et al.* (1995); Natarajan (1995). An alternate approach is to formulate a convex relaxation of this NP-hard problem by introducing slack variables into the constraints in (27), and then minimizing an appropriate norm of the vector of slack variables. If we choose a particular norm $\|\cdot\|$ to measure distances in "feature space", then distances in "weight space" should be measured using the so-called **dual norm**, defined by

$$\|w\|_d := \sup_{\|x\| \leq 1} |xw|.$$

In particular, if we measure distances in feature space using the ℓ_1-norm,then distances in weight space should be measured using its dual, which is the ℓ_∞-norm. With this observation, the problem can be formulated as follows:

$$\min_{w,\theta,y,z} (1-\lambda) \left[\sum_{i=1}^{m_1} y_i + \sum_{i=1}^{m_2} z_i \right] + \lambda \max_{1 \leq i \leq n} |w_i| \text{ s.t.}$$

$$x^i w - \theta + y_i \geq 1 \; \forall i \in \mathcal{M}_1, x^i w - \theta - z_i \leq -1 \; \forall i \in \mathcal{M}_2,$$

$$y \geq \mathbf{0}_{m_1}, z \geq \mathbf{0}_{m_2}. \tag{28}$$

This can be converted to

$$\min_{w,\theta,y,z} (1-\lambda) \left[\sum_{i=1}^{m_1} y_i + \sum_{i=1}^{m_2} z_i \right] + \lambda v \text{ s.t.}$$

$$x^i w - \theta + y_i \geq 1 \; \forall i \in \mathcal{M}_1, x^i w - \theta - z_i \leq -1 \; \forall i \in \mathcal{M}_2,$$

$$y \geq \mathbf{0}_{m_1}, z \geq \mathbf{0}_{m_2}, v \geq w_i \; \forall i, v \geq -w_i \; \forall i. \tag{29}$$

This is clearly a linear programming problem. In this formulation, λ is a "small" constant in $(0, 1)$, much closer to 0 than it is to 1. Suppose that the original data set is linearly separable, and let w^*, θ^* denote a solution to the optimization problem in (27), where $\|w\|_d$ replaces $\|w\|$. Then the choice

$$w = w^*, \theta = \theta^*, y = \mathbf{0}_{m_1}, z = \mathbf{0}_{m_2}$$

is certainly *feasible* for the optimization problem (28). Moreover, if λ is sufficiently small, any reduction in $\|w\|_d$ achieved by violating the linear separation constraints (i.e., permitting some y_i or z_i to be positive rather than zero) is offset by the increase in the term $(1-\lambda)\|(y, z)\|$. It is therefore clear that, if the data set is linearly separable, then there exists a critical

value $\lambda_0 > 0$ such that, for all $\lambda < \lambda_0$, the optimization problem (28) has $(w^*, \theta, \mathbf{0}_{m_1}, \mathbf{0}_{m_2})$ as a solution. On the other hand, the optimization problem (28) remains meaningful even when the data is not linearly separable.

The final aspect of the problem, as suggested in Veropoulos *et al.* (1999), is to introduce a trade-off between false positives and false negatives. In this connection, it is worthwhile to recall the definitions of the accuracy etc. of a classifier. Given a discriminant function $f(\cdot)$, define

$$\mathcal{C}_1 := \{i \in \mathcal{M} : f(x^i) > 0\}, \mathcal{C}_2 := \{i \in \mathcal{M} : f(x^i) < 0\}.$$

Thus $\mathcal{C}_1$ consists of the samples that are assigned to Class 1 by the classifier, while $\mathcal{C}_2$ consists of the samples that are assigned to Class 2. Then this leads to the array shown below:

$$\begin{array}{ccc} & \mathcal{C}_1 & \mathcal{C}_2 \\ \mathcal{M}_1 & TP & FN \\ \mathcal{M}_2 & FP & TN \end{array}$$

In the above array, the entries TP, FN, FP, TN stand for "true positive", " false negative", "false positive" and "true negative" respectively.

Definition 4.1 *With the above definitions, we have*

$$Se = \frac{TP}{TP + FN} = \frac{|\mathcal{C}_1 \cap \mathcal{M}_1|}{|\mathcal{M}_1|}, \tag{30}$$

$$Sp = \frac{TN}{FP + TN} = \frac{|\mathcal{C}_2 \cap \mathcal{M}_2|}{|\mathcal{M}_2|}, \tag{31}$$

$$Ac = \frac{TP + TN}{TP + TN + FP + FN} = \frac{|\mathcal{C}_1 \cap \mathcal{M}_1| + |\mathcal{C}_2 \cap \mathcal{M}_2|}{|\mathcal{M}_1| + |\mathcal{M}_2|}, \tag{32}$$

where Se, Sp, Ac stand for the **sensitivity**, **specificity**, *and* **accuracy** *respectively.*

All three quantities lie in the interval $[0, 1]$. Moreover, accuracy is a convex combination of sensitivity and specificity. In particular,

$$Ac = Se \cdot \frac{|\mathcal{M}_1|}{|\mathcal{M}_1| + |\mathcal{M}_2|} + Sp \cdot \frac{|\mathcal{M}_2|}{|\mathcal{M}_1| + |\mathcal{M}_2|}.$$

Therefore

$$\min\{Se, Sp\} \leq Ac \leq \max\{Se, Sp\}.$$

Also, the accuracy of a classifier will be roughly equal to the sensitivity if $\mathcal{M}_1$ is far larger than $\mathcal{M}_2$, and roughly equal to the specificity if $\mathcal{M}_2$ is far larger than $\mathcal{M}_1$.

In many classification problems the consequences of misclassification are not symmetric. To capture these kinds of considerations, another parameter $\alpha \in (0, 1)$ is introduced, and the objective function in the optimization problem (28) is modified by making the substitution

$$\sum_{i=1}^{m_1} y_i + \sum_{i=1}^{m_2} z_i \leftarrow \alpha \sum_{i=1}^{m_1} y_i + (1 - \alpha) \sum_{i=1}^{m_2} z_i,$$

where we adopt the computer science notation $\leftarrow$ to mean "replaces." If $\alpha = 0.5$, then both false positives and false negatives are weighted equally. If $\alpha > 0.5$, then there is greater emphasis on correctly classifying the vectors in $\mathcal{M}_1$, and the reverse if $\alpha < 0.5$. With this final problem formulation, the following desirable properties result:

- The problem is a linear programming problem and is therefore tractable for even for extremely large values of n, the number of features.
- The formulation can be applied without knowing beforehand whether or not the data set is linearly separable.
- The formulation provides for a trade-off between false positives and false negatives.
- Most important, the optimal weight vector w has at most m nonzero entries, where m is the number of samples. Hence the classifier uses at most m out of the n features.

For these reasons, the ℓ_1-norm SVM forms the starting point for our further research into classification.

4.3 The Lone Star Algorithm

As pointed out in Section 4, both the traditional ℓ_2-norm support vector machine (SVM) as well as the ℓ_1-norm SVM can be used for two-class classification problems. When the number of samples m is far larger than the number of features n, the traditional SVM performs very satisfactorily, whereas the ℓ_1-norm SVM of Bradley and Mangasarian (1998) is to be preferred when $m < n$. Moreover, the ℓ_1-norm SVM is guaranteed to use no more than m features. However, in many biological applications, even m features are too

many. Biological measurements suffer from poor repeatability. Therefore a classifier that uses fewer features would be far preferable to one that uses more features. In this section we present a new algorithm for two-class classification that often uses far fewer than m features, thus making it very suitable for biological applications. The algorithm combines the ℓ-norm SVM of Bradley and Mangasarian (1998), recursive feature elimination (RFE) of Guyon *et al.* (2002), and stability selection of Meinshausen and Bühlmann (2010). A preliminary version of this algorithm was reported in Ahsen *et al.* (2012).

The algorithm is as follows:

1. Choose at random a "training set" of samples of size k_1 from $\mathcal{M}_1$ and size k_2 from $\mathcal{M}_2$, such that $k_l \leq m_l/2$, and k_1, k_2 are roughly equal. Repeat this choice s times, where s is a "large" number. This generates s different "training sets", each of which consists of k_l samples from $\mathcal{M}_l$, $l = 1, 2$.

2. For each randomly chosen training set, compute a corresponding optimal ℓ_1-norm SVM using the formulation (28). This results in s different optimal weight vectors and thresholds.

3. Let k denote the average number of nonzero entries in the optimal weight vector across all randomized runs. Average all s optimal weight vectors and thresholds, retain the largest k components of the averaged weight vector and corresponding feature set, and set the remaining components to zero. This results in reducing the number of features from the original n to k.

4. Repeat the process with the reduced feature set, but the originally chosen randomly selected training samples, until no further reduction is possible in the number of features. This determines the final set of features to be used.

5. Once the final feature set is determined, carry out two-fold cross validation by dividing the data s times into a training set of k_1, k_2 randomly selected samples and assessing the performance of the resulting ℓ_1-norm classifier on the testing data set, which is the remainder of the samples. Average the weights generated by the $t \leq s$ best-performing classifiers where t is chosen by the user, and call that the final classifier.

When the number of features n is extremely large, an optional preprocessing step is to compute the mean value of each of the n features for

each class, and retain only those features wherein the difference between means is statistically significant using the "Student" t-test. Our experience is that using this optional pre-processing step does not change the final answer very much, but does decrease the CPU time substantially.

Now some comments are in order regarding the above algorithm.

- In some applications, $\mathcal{M}_1$ and $\mathcal{M}_2$ are of comparable size, so that the size of the training set can be chosen to equal roughly half of the total samples within each class. However, in other applications, the sizes of the two sets are dissimilar, in which case the larger set has far fewer of its samples used in training.

- Step 1 of randomly choosing s different training sets differs from Guyon *et al.* (2002), where there is only one randomized division of the data into training and testing sets.

- For each random choice of the training set, the *number* of nonzero entries in the optimal weight vector is more or less the same; however, the *locations* of nonzero entries in the optimal weight vector vary from one run to another.

- In Step 3 above, instead of averaging the optimal weights over all s runs and then retaining the k largest components, it is possible to adopt another strategy. Rank all n indices in order of the number of times that index has a nonzero weight in the s randomized runs, and retain the top k indices. In our experience, both approaches lead to virtually the same choice of the indices to be retained for the next iteration.

- Instead of choosing s randomized training sets right at the outset, it is possible to choose s randomized training sets each time the number of features is reduced.

- In the final step, there is no distinction between the training and testing data sets, so the final classifier is run on the entire data set to arrive at the final accuracy, sensitivity and specificity figures.

The advantage of the above approach vis-a-vis the ℓ_2-norm SVM-RFE of Guyon *et al.* (2002) is that the number of features reduces significantly at each step, and the algorithm converges in just a few steps. This is because, in the ℓ_1-norm SVM, many components of the weight vector are "naturally" zero, and need not be truncated. In contrast, in general all the components

of the weight vector resulting from the ℓ_2-norm SVM will be nonzero; as a result the features can only be eliminated one at a time, and in general the number of iterations is equal to (or comparable to) n, the initial number of features.

The new algorithm can be appropriately referred to as the "ℓ_1-SVM t-test and RFE" algorithm, where SVM and RFE are themselves acronyms as defined above. Once again taking the first letters, we are led to the "second-level" acronym "ℓ_1-StaR", which can be pronounced as "ell-one star". Out of deference to our domicile, we have decided to call it the "lone star" algorithm.

The lone star algorithm was applied to the problem of predicting which patients of endometrial cancer are at risk of lymph node metastasis. These results are reported elsewhere. But in brief the situation is the following: The endometrium is the lining of the uterus. When a patient contracts endometrial cancer, her uterus, ovaries, and fallopian tubes are surgically removed. One of the major risks run by endometrial cancer patients is that the cancer will metastasize and spread through the body via pelvic and/or para-aortic lymph nodes. The Gynecological Oncology Group (GOG) recommends that the patient's pelvic and para-aortic lymph nodes should also be surgically removed when the size of the tumor exceeds 2cm in diameter. However, post-surgery analysis reveals that even in this case, lymphatic metastasis is present in only 22% of the cases Mariani *et al.* (2008).

To predict the possibility of lymphatic metastasis, 1,428 micro-RNAs were extracted from 94 tumors, half with and half without metastasis. Using the lone star algorithm, 13 micro-RNAs were identified as being highly predictive. When tested on the entire training sample of 94 tumors, the lone star classifier correctly classified 41 out of 43 lymph-positive samples, and 40 out of 43 lymph-negative samples. In on-going work, these micro-RNAs were measured on an independent cohort of 19 lymph-negative and 9 lymph-positive tumors. The classifier classified 8 out of 9 lymph-positive tumors correctly, and 11 out of 19 lymph-negative tumors correctly. Thus, while the specificity is not very impressive, the sensitivity is extremely good, which is precisely what one wants in such a situation. Moreover, using a two-table contingency analysis and the Barnard exact test, the likelihood of arriving at this assignment by pure chance (the so-called P-value) is bounded by 0.011574. In biology any P-value less than 0.05 is generally considered to be significant.

5 Some Topics for Further Research

Machine learning and computational biology are both vast subjects, and their intersection contains many more topics than are touched upon in this brief article. Besides, there are other topics in computational cancer biology that do not naturally belong to machine learning, for example, modelling tumor growth using branching processes. Therefore the emphasis in this article has been on topics that are well-established in the machine learning community, and are also relevant to problems in computational cancer biology.

Until now several "penalty" norms have been proposed for inducing an optimization algorithm to select structured sparse feature sets, such as group lasso (GL) and sparse group lasso (SGL). As pointed out in Section 2, available extensions of these penalty norms to overlapping sets do not address biological networks where there are multiple paths from a master regulator to a final node. Any advance in this direction would have an immediate application to computational biology.

Compressed sensing theory as discussed in Section 3 is based on the premise that *is possible to choose* the measurement matrix A. The available theorems in this theory are based on assumptions on the measurement matrix, such as the restricted isometry property, or the null space property, and perhaps something even more general in future. In order to apply techniques from compressed sensing theory to cancer biology, it would be necessary to modify the theory to the case where the measurement matrix is given, and not chosen by the user. The RIP corresponds to the assumption that in an $m \times n$ matrix A, every choice of k columns results in a nearly orthogonal set. In actual biological data, such an assumption has no hope of being true, because the expression levels of some genes would be highly correlated with those of other genes. In Candès and Plan (2009), the authors suggest that it is possible to handle this situation by first clustering the column vectors and then choosing just one exemplar from each cluster before applying the theory. Our preliminary attempts to apply such an approach to ovarian cancer data (The Cancer Genome Atlas Network (2011)) are not very promising, leading to RIP orders of 5 or 10 – far too small to be of practical use. Thus there is a need for the development of other heuristics besides clustering to extract nearly orthogonal sets of columns for actual measurement matrices. In this connection it is worth pointing out Huang and Zhang (2010) that group RIP is easier to achieve using random projections, as compared to RIP. However, it is not clear whether a "given" A matrix is likely to satisfy a group RIP with a sufficiently large order.

In general it would appear that sparse regression is more advanced than sparse classification, with both well-established theoretical foundations as well as widely used algorithms in the former. In contrast, sparse classification does not have such a wealth of results. The lone star algorithm introduced here has performed well in several applications involving cancer data, and at least for the moment, it appears to be the only available method to select far fewer features than the size of the training set of samples. As of now there is no theoretical justification for this observed behavior. Recall that the ℓ_1-norm SVM is guaranteed only to choose no more features than the size of the training set; but there is no reason to assume that it will use fewer. Therefore it is certainly worthwhile to study when and why lone star and other such algorithms will prove to be effective.

Acknowledgement

The author would like to assert his deep gratitude to Professor Michael A. White of the UT Southwestern Medical Center, Dallas, TX for introducing him to the fascinating world of cancer biology and for turning him into a passable imitation of a computational cancer biologist. He would also like to thank his students Eren Ahsen, Burook Misganaw and Nitin Singh for carrying out much of the work that supports the theory reported here.

References

Achlioptas, D. (2003). Database-friendly random projections: Johnson-Lindenstraus with binary coins. *Journal of Computer and System Sciences*, **66**, 671–687.

Ahsen, M. E. and Vidyasagar, M. (2014). Near-ideal behavior of compressed sensing algorithms. *arXiv*, page 1401.6623.

Ahsen, M. E., Singh, N. K., Boren, T., Vidyasagar, M., and White, M. A. (2012). A new feature selection algorithm for two-class classification problems and application to endometrial cancer. In *Proceedings, IEEE Conference on Decision and Control, Maui, Hawaii*.

Bakin, S. (1999). Adaptive regression and model selection in data mining problems. Ph.D. thesis, The Australian National University, Canberra.

Baraniuk, R., Davenport, M., Devore, R., and Wakin, M. (2008). A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, **28**, 253–263.

Bradley, P. S. and Mangasarian, O. L. (1998). Feature selection via concave minimization and support vector machines. In *Machine Learning: Proceedings of the Fifteenth International Conference (ICML '98)*, pages 82–90. Morgan Kaufmann, San Francisco.

Cancer Research UK (2013). http://www.cancerresearchuk.org.

Candès, E. (2008). The restricted isometry property and its implications for compresed sensing. *Comptes rendus de l'Académie des Sciences, Série I*, **346**, 589–592.

Candès, E. J. and Plan, Y. (2009). Near ideal model selection by ℓ_1 minimization. *The Annals of Statistics*, **37(5A)**, 2145–2177.

Candès, E. J. and Tao, T. (2005). Decoding by linear programming. *IEEE Transactions on Information Theory*, **51**, 4203–4215.

Candès, E. J. and Tao, T. (December 2007). The Dantzig selector: Statistical estimation when p is much larger than n. *Annals of Statistics*, **35(6)**, 2313–2351.

Cohen, A., Wolfgang, Dahmen, and Devore, R. (January 2009). Compressed sensing and best k-term approximation. *Journal of the American Mathematical Society*, **22(1)**, 211–231.

Cortes, C. and Vapnik, V. N. (1997). Support vector networks. *Machine Learning*, **20**.

COSMIC (2013). http://www.sanger.ac.uk/genetics/CGP/cosmic.

Cristianini, N. and Shawe-Taylor, J. (2000). *Support Vector Machines*. Cambridge University Press, Cambridge, UK.

Davenport, M. A., Duarte, M. F., Eldar, Y. C., and Kutyniok, G. (2012). Introduction to compressed sensing. In Y. C. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*, pages 1–68. Cambridge University Press.

Donoho, D. (2006a). Compressed sensing. *IEEE Transactions on Information Theory*, **52(4)**, 1289–1306.

Donoho, D. (2006b). For most large underdetermined systems of linear equations, the minimal ℓ_1-norm solution is also the sparsest solution. *Communications in Pure and Applied Mathematics*, **59(6)**, 797–829.

Friedman, J., Hastie, T., and Tibshirani, R. (2010). A note on the group lasso and sparse group lasso. http://www-stat.stanford.edu/ tibs/ftp/sparse-grlasso.pdf.

GEO (2013). http://www.ncbi.nlm.nih.gov/geo/.

Guyon, I., Weston, J., Barnhill, S., and Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. *Machine Learning*, **46**, 389–422.

Hastie, T., Tibshirani, R., and Friedman, J. (2011). *The Elements of Statistical Learning (Second Edition)*. Springer, New York.

Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, **12(1)**, 55–67.

Höffgen, K.-U., Simon, H.-U., and Horn, K. S. V. (1995). Robust trainability of single neurons. *Journal of Computer and System Science*, **50(1)**, 114–125.

Huang, J. and Zhang, T. (2010). The benefit of group sparsity. *The Annals of Statistics*, **38(4)**, 1978–2004.

International Cancer Genomics Consortium (2013). http://icgc.org/.

Jenetton, R., Mairal, J., Obozinski, G., and Bach, F. (2011). Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, **12**, 2297–2334.

Lin, Y. and Zhang, H. (2006). Component selection and smoothing in smoothing spline analysis of variance models. *Annals of Statistics*, **34**, 2272–2297.

Mariani, A., Dowdy, S. C., Cliby, W. A., Gostout, B. S., Jones, M. B., Wilson, T. O., *et al.* (2008). Prospective assessment of lymphatic dissemination in endometrial cancer: a paradigm shift in surgical staging. *Gynecol Oncol.*, **109(1)**, 11–18.

Meinshausen, N. and Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society, Part B*, pages 417–483.

Natarajan, B. K. (1995). Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, **24**, 227–234.

Negabhan, S., Ravikumar, P., Wainwright, M. J., and Yu, B. (December 2012). A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Statistical Science*, **27(4)**, 538–557.

Northrop, R. B. and Connor, A. N. (2009). *Introduction to Molecular Biology, Genomics and Proteomics for Biomedical Engineers*. CRC Press.

Obozinski, G., Jacob, L., and Vert, J.-P. (2011). Group lasso with overlaps: The latest group lasso approach. *arxiv*, page 1110.0413.

Osborne, M. R., Presnell, B., and Turlach, B. A. (2000). On the LASSO and its dual. *Journal of Computational and Graphical Statistics*, **9**, 319–337.

Romberg, J. (2009). Compressive sensing by random convolution. *SIAM Journal on Imaging Sciences*, **2(4)**, 1098–1128.

Schölkopf, B. and Smola, A. J. (2002). *Learning with Kernels*. MIT Press.

SEER (2013). http://seer.cancer.gov/statfacts/html/all.html.

Siegel, R., Naishadham, D., and Jemal, A. (January/February 2013). Cancer statistics, 2013. *CA: A Cancer Journal for Clinicians*, **63(1)**, 11–30.

Simon, N., Friedman, J., Hastie, T., and Tibshirani, R. (2012). A sparse group lasso. http://www-stat.stanford.edu/ nsimon/SGLpaper.pdf.

The Cancer Genome Atlas (2013). http://cancergenome.nih.gov/.

The Cancer Genome Atlas Network (30 June 2011). Integrated genomic analyses of ovarian carcinoma. *Nature*, **474**, 609–615.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, **58(1)**.

Tikhonov, A. N. (1943). On the stability of inverse problems. *Doklady Akademii Nauk SSSR*, **39(5)**, 195–198.

Tözeren, A. and Byers, S. W. (2003). *New Biology for Engineers and Computer Scientists*. Prentice-Hall.

Veropoulos, K., Campbell, C., and Cristianini, N. (1999). Controlling the sensitivity of support vector machines. In *IJCAI Workshop on Support Vector Machines*.

Vidyasagar, M. (2003). *Learning and Generalization: With Applications to Neural Networks and Control Systems*. Springer-Verlag, London.

Wenocur, R. S. and Dudley, R. M. (1981). Some special Vapnik-Chervonenkis classes. *Discrete Mathematics*, **33**, 313–318.

World Health Organization (2013). http://www.who.int/cancer/en/.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. Royal Stat. Soc. B*, **67**, 301–320.